

Knowledge Management in High-Hazard Industries

Accident Precursors as Practice

JOHN S. CARROLL
Sloan School of Management
Massachusetts Institute of Technology

Accident precursors are events that must occur for an accident to happen in a given scenario, but that have not resulted in an accident so far. High-hazard industries, such as nuclear power and aviation, that would put many people at risk in the event of a single accident are particularly sensitive to precursors and consider them opportunities to avoid accidents. Accidents happen when precursors occur in combination and/or when system defenses fail to mitigate a situation. Every precursor event is, therefore, both a test of the adequacy of system defenses and an opportunity to develop and apply knowledge to avoid accidents. Failure to take notice of these “tests” and to build a strong knowledge-management system is a sign of trouble ahead.

At the Three Mile Island (TMI) nuclear power plant, for example, a combination of events—a stuck-open pressure-relief valve that allowed water levels in the reactor to drop, thus uncovering the radioactive core plus indicators that showed the position of the switch controlling the valve but not the valve itself plus operator training that cautioned operators about overfilling the reactor with water—destroyed a billion dollar unit of the plant and changed the nuclear power industry forever. Even though information that could have prevented the TMI event was available from similar incidents at other plants, recurrent problems with the same equipment at TMI, and critiques of operator training, that information was not incorporated into plant-wide or industry-wide operating practices (Marcus et al., 1989). The president of the utility, Herman Dieckamp later reflected on the incident (Kemeny et al., 1979):

To me that is probably one of the most significant learnings of the whole accident, the degree to which the inadequacies of that experience feedback loop . . . significantly contributed to making us and the plant vulnerable to this accident.

In response to the TMI accident, the nuclear power industry created the Institute of Nuclear Power Operations (INPO) to identify precursors, disseminate lessons learned and best practices, and generally ensure that every plant operates with the best knowledge available (and also to forestall further regulation). The World Association of Nuclear Operators performs these tasks globally. Although knowledge development and dissemination have been successful overall, problems continue in this industry, which is under continuous scrutiny by regulators and a wary public.

ACCIDENT PRECURSORS AND KNOWLEDGE MANAGEMENT

From a knowledge-management perspective, precursors are signals of possible problems, chinks in an operation's armor, or pathways to accidents. They are called precursors rather than accidents because systems have multiple layers of defense like slices of Swiss cheese stacked together (Reason, 1997). A precursor problem may pass through one or two layers of defense (through the holes in the Swiss cheese), but another layer usually stops the progression toward an accident. Only when "all of the holes line up" does the problem overcome or bypass all defenses and become an accident. As signals, precursors allow us to find the sources of potential problems and assess the robustness of defenses. Based on information from precursors, we can improve defenses or make sure they function as designed and add new defenses when problems become frequent or serious or new problems appear.

The history of the nuclear power industry shows a constant tension between wariness and complacency. Early on, operators and regulators believed that nuclear power would be a simple technology to operate, that electricity would be "too cheap to meter," and that safety would be assured. TMI was a "fundamental surprise" (Lanir, 1986) that caused intense scrutiny and huge investments in safety equipment, procedures, training, reporting, and people. Probabilistic risk analysis was invented as a way of anticipating problems and designing defenses against them. However, each time the industry has thought its was secure in its ability to anticipate problems and design defenses, new, unanticipated challenges have arisen, such as shut-down risk, stress corrosion cracking, and inadequacies in safety culture. The industry continues to learn, forget, and relearn a difficult lesson—that anticipation must be combined with resilience in responding to precursors (Marcus and Nichols, 1999; Weick et al., 1999; Wildavsky, 1988).

One institutionalized approach to combating problems and remaining alert is self-assessment embedded in corrective-action programs. In a speech to the Regulatory Information Conference in 1996, Dr. Shirley Jackson, former chair

of the U.S. Nuclear Regulatory Commission (USNRC) attributed improvement in the 1990s to “increased emphasis by both the [US] NRC and the industry in the following three areas: (1) improved maintenance practices; (2) consideration of risk in the operation and maintenance of nuclear plants; and (3) self-assessment of events to identify root causes of problems and ensure effective corrective actions.” She went on to say that self-assessment “should be an ingrained part of every licensee’s way of doing business” and that self-assessment would become increasingly important as the industry moved “to more performance-oriented regulatory approaches.”

THE STOCK-AND-FLOW MODEL OF KNOWLEDGE MANAGEMENT

Traditional knowledge management is a combination of maintaining repositories of explicit information and expert know-how organized by professional discipline. Examples of explicit information include databases, procedural manuals, drawings, and planning documents. Routine operations are guided by this codified knowledge, and routine problems can often be addressed by consulting the manuals. Thus, some knowledge can be explicitly codified in these reservoirs (Argote and Ingram, 2000), but some knowledge is tacit, implicit in the experience and training of individuals. Thus, engineers, operators, craftsmen, accountants, and others with expertise in particular domains have developed “judgment” and recognition-based diagnostic and action skills (Klein, 1998). Most exceptions and problems can be categorized and referred to subject-matter experts for resolution.

In this model of knowledge management, the key issue is “where” the knowledge resides. Knowledge is a stock or supply that has to be accessible and can be moved around as needed, like supplies in a warehouse or money in a bank account. When a precursor is noticed, a search is made for relevant information to ensure that defenses are adequate or to strengthen defenses if necessary. The search focuses on the problem (e.g., if the problem has been seen before, if other plants in the industry have seen it) and on the domains of expertise relevant to the problem (e.g., maintenance, engineering, chemistry). Investigators have access to databases created by a plant, groups of similar plants, manufacturers, industry groups, and even regulators. Explicit knowledge in the databases can be applied directly, and deviations are dealt with by evolutionary enhancements, including adding controls: “Safe operating procedures . . . are continually being amended to prohibit actions that have been implicated in some recent accident or incident” (Reason, 1997).

However, most problems involve knowledge that is local and contextual, tacit as well as explicit. Therefore, additional knowledge is necessary before what is known can be applied to new instances. In other words, problems may not be identical from place to place or time to time, and information may be

“sticky” or difficult to move from one location to another (Szulanski, 1996; von Hippel, 1994). Expert judgment may be necessary to draw analogies, tailor solutions for particular situations, and so forth. In such cases, success depends upon the personal involvement of knowledgeable individuals and personal networks that connect accountable investigators with knowledgeable experts.

Industries such as nuclear power recognize the importance of personal contacts in the dissemination of best practices, experience with precursors, and so forth. Virtually all bits of new information include contact information for individuals who are the best sources of information. Thus, the article or the database entry is an advertisement or infomercial rather than a source of necessary information. To implement a best practice, one must learn by telephone, by visiting the source plant, by hosting peer-assist visits from source-plant personnel, or by using consultants as transmission channels. Contacts may be facilitated by liaisons, job rotations, or temporary exchanges of personnel with other plants or industry organizations, such as INPO. Thus, knowledge management depends upon the development of informal (often invisible) networks of personal contacts within a plant, with other plants, with suppliers, consultants, regulators, universities, etc. One of the first cultural precursors to trouble is an organization that withdraws from “nonessential” industry activities and, therefore, limits its access to new information and knowledgeable peers; this is what Millstone Station did in the 1980s following the financial challenges of building a third unit (Carroll and Hatakenaka, 2001).

Hansen (1999) showed that different kinds of network ties or interpersonal relationships are necessary for different kinds of knowledge transfer. Having a large number of “weak ties,” that is, infrequent, distant relationships and acquaintanceships, facilitates the search for new knowledge. A person with a broad network can find new information easily, including by using e-mail and web searches. If the information is relatively simple and easy to transfer, weak ties are very efficient and useful. However, weak ties can actually slow down the transmission of complex information, which requires a strong connection among individuals or groups.

THE CAPABILITIES MODEL OF KNOWLEDGE MANAGEMENT

We can conceptualize knowledge management as a system capable of attending to signals, generating new knowledge (updating), retaining knowledge, and applying knowledge where it is needed. This constellation of capabilities is sometimes called organizational learning (Carroll et al., 2002; Crossan et al., 1999; Senge, 1990). For our purposes, organizational learning is another description of how knowledge is generated and applied in action, which includes capabilities for attending, making sense, and implementing change.

Attention or “heedfulness” is a crucial first step in reacting to precursors (Marcus and Nichols, 1999; Weick et al., 1999). In most organizations, precursors

either go unheeded or are responded to at the local level with no signal reaching beyond the immediate work context. Reporting systems are an institutionalized form of attention; planning, typically understood as a way of allocating resources and controlling activities, enables people within an organization to notice things more easily and to get more rapid and more useful feedback about how things are going (deGeus, 1988).

Organizations rarely succeed because they “meet plan,” but organizations without a clear plan find it hard to notice when things are not going well and, therefore, to respond to incipient problems creatively and effectively. For precursors to be recognized as precursors, there has to be a shared understanding of what is normal and what is off-normal, what is expected and what is unexpected, what is desirable and what is undesirable. As Weick and Sutcliffe (2001) state, “to move toward high reliability is to enlarge what people monitor, expect, and fear.” A typical nuclear power plant, for example, formally identifies more than 2,000 problems or incidents per year, 90 percent of which would have been ignored a decade ago.

Once precursors or troublesome conditions have been noticed, some type of analysis or investigation follows. Nearly all high-hazard organizations conduct investigations of problems as part of their corrective-action programs, which start with the reporting of problems and continue with the investigation of facts and opinions, the attribution of causes, the generation of insights and recommendations, the implementation of interventions to improve performance, and the verification that these interventions are carried out and produce the expected results (Carroll, 1995, 1998; Carroll et al., 2001; van der Schaaf et al., 1991). More frequent than the massive investigations triggered by rare accidents, such as TMI, these smaller scale self-analyses and problem-solving activities focus on small defects, near misses, and other lesser failures (Sitkin, 1992) or precursors (Reason, 1990). Problem investigation is a kind of off-line, reflective practice that involves sense-making, analysis, and imagining alternatives. This often takes place outside of the regular work process, often by individuals who are not immediately involved in the problem (Argyris, 1996; Rudolph et al., 2001).

Although individuals can investigate most problems, the most serious, persistent, causally ambiguous, and organizationally complex problems are investigated by teams. Each year, nuclear power plants assemble multidisciplinary teams (sometimes including personnel from other plants, headquarters, other companies, and elsewhere) to investigate a small number of problems that seem to extend beyond the knowledge base of any single department. These teams not only provide a wide range of expertise, they also have better access to information from informants and more credibility with the audiences who must support the implementation of their recommendations. Serving on these teams can provide valuable experience and enhance an individual’s knowledge and skills, which are then brought back to coworkers when the team member returns to his or her home department (Gruenfeld et al., 2000); the experience helps bridge the

gap between communities of practice, thus enhancing the capabilities of the organization as a whole (Cook and Brown, 1999).

Investigations often focus on fixing immediate problems so operations can return to normal and everyone can regain a sense of predictability and control, which are so important to managers and engineers, especially in high-hazard industries (Carroll, 1998; Carroll et al., 2002). However, just as exploiting readily available information may keep one from exploring new possibilities (March, 1991), fixing immediate problems may interfere with the extraction of all useful information from a precursor event. For example, in the chemical plant pipe failure reported by Hendershot et al. (2003) or the chemical plant charge-heater fire reported by Carroll et al. (2002), investigations could have stopped with simple explanations and fixes that would have prevented those particular problems from recurring. In both cases, however, the analyses went further to identify "root causes," which resulted in new knowledge about the technology and organization of the work.

In the charge-heater fire investigation, for example, the team noted as a "Key Learning" that the plant staff had made decisions without questioning their assumptions. First, the maintenance department had changed decoking processes but did not know and never checked to be sure that the new process was effective. Second, operators increased the burner pressure in the charge heater to increase production but did not know the consequences of doing so. Third, operators changed the pattern of firing heater tubes (to fire hotter around the perimeter without setting off safety alarms) but again did not know the consequences of doing so. The investigation team found that the fire was caused by a combination of (1) operators firing heater tubes in such a way that the hottest temperatures were located away from the instruments designed to detect danger and (2) the presence of residual coke (coal dust) on the inside of the tubes that the new maintenance process had left behind. On the basis of these findings, the first recommendation for future action was that the plant identify "side effects" and be more aware of the broad "decision context" when changing production processes. This resulted in the implementation of a new "management of change" process so that the global implications of proposed local actions could be anticipated better.

THE PRACTICE MODEL OF KNOWLEDGE MANAGEMENT

Neither the stock-and-flow model nor the capabilities model describes in detail how knowledge management is accomplished. The assumption is that the right tools, people, and environment will promote the development, transfer, and use of knowledge. The practice model of knowledge management focuses on specific activities (Bourdieu, 1977; Brown and Duguid, 1991; Carlile, 2002). For example, knowledge is often embodied in stories and transmitted through storytelling. In addition, knowledge development among communities-of-practice requires

specific boundary-spanning or bridging practices. Incident investigations and analyses of root causes (which include a variety of techniques for looking beyond immediate or proximal causes) may be valuable not only as analytical tools, but also as opportunities for conversations with shared purposes (Carroll et al., 2002).

In our research on incident investigation teams in nuclear power plants, we assumed that teams that used root-cause analysis to make deeper investigations of precursor events would generate more knowledge and that organizations would implement more effective changes that would improve performance. We discovered, however, that the investigation teams and the managers to whom they reported had very different ideas about what constituted a good investigation and a good report. The teams wanted to find the causes of precursors, to dig deeply and identify failed defenses. The managers wanted actionable recommendations that would fix problems and reestablish control. Managers seeking efficiency delegated participation on the team and waited to respond to a draft report rather than taking the time to work directly with the team (Nutt, 1999). As a result, the hand-off from team to manager was often ineffective. Reports were sometimes modified or negotiated to obtain manager "sign off," and recommendations were sometimes watered down or folded into other activities, or even refused, on the basis of cost or other practicalities. Managers often thought investigation teams were unrealistic, whereas the teams thought managers were defensive.

Interestingly, at the chemical company that investigated the charge-heater fire, the investigation team had an explicit goal of *educating managers*, rather than solving problems! In this company, teams presented facts and carefully reasoned causal connections, but did not make recommendations. The managers' collective job was to understand the problem and its context, discuss opportunities for improvement, commission activities to develop solutions, and implement changes.

Problem investigations provide precisely the kind of opportunities that can bring together diverse perspectives and facilitate learning and change. The mixing of occupational and educational backgrounds (Dougherty, 1992; Rochlin and von Meier, 1994) and cognitive styles (Jackson, 1992; White, 1984) that combine abstract, systemic issues with concrete, operational details and technical complexity with human ambiguity can lead to informational diversity (Jehn et al., 1999) or "conceptual slack" (Schulman, 1993). Weick et al. (1999) similarly argue that consistent reliability requires that problems not be oversimplified, which requires diverse perspectives and frequent boundary-spanning activities. This creates skills and opportunities for engaging in a process of knowing that can bring to the surface previously unarticulated mental models of the work environment, compare them, and lead to new, shared models (Cook and Brown, 1999).

In the cases we studied, boundary spanning was only partially successful. Sharpening and bridging differences among disciplines and hierarchical levels

requires an atmosphere of mutual respect and trust. Managers, however, were often not full participants on the investigative teams, reports were rather casual in connecting causes and recommended actions, and negotiations over the final report tended to be about authority rather than reasoning. It takes mindful attention to build shared understanding around diffuse issues, such as “culture” and “accountability,” that have very different meanings and implications to different professional groups (Carroll, 1998; Carroll et al., 2002). Because the emphasis is usually on controlling deviations from existing procedures and rules, few teams and managers are willing or able to work hard to clarify meaning and build shared mental models. Therefore, they often miss opportunities to deepen their understanding that could lead to organizational learning and change.

SUMMARY

All politics is said to be local, and in important ways knowledge is local as well. In managing knowledge about accident precursors, organizations must attend to the local nature of problems and the knowledge that must be brought to bear to address them, as well as to the global nature of what is learned and what may be needed at other times in other locations. In addition, they must consider knowledge not only as a stock of information, but also as providing the capability of inquiring, imagining, bridging boundaries, building networks of trusting relationships, and taking action. Precursor events are opportunities to enact and improve organizational practices.

REFERENCES

- Argote, L., and P. Ingram. 2000. Knowledge transfer: a basis for competitive advantage in firms. *Organizational Behavior and Human Decision Processes* 82: 150–169.
- Argyris, C. 1996. Unrecognized defenses of scholars: impact on theory and research. *Organization Science* 7: 79–87.
- Bourdieu, P. 1977. *Outline of a Theory of Practice*. Cambridge, U.K.: Cambridge University Press.
- Brown, J.S., and P. Duguid. 1991. Organizational learning and communities-of-practice: toward a unified view of working, learning, and innovation. *Organization Science* 2: 40–57.
- Carlile, P.R. 2002. A pragmatic view of knowledge and boundaries: boundary objects in new product development. *Organization Science* 13: 442–455.
- Carroll, J.S. 1995. Incident reviews in high-hazard industries: sense-making and learning under ambiguity and accountability. *Industrial and Environmental Crisis Quarterly* 9: 175–197.
- Carroll, J.S. 1998. Organizational learning activities in high-hazard industries: the logics underlying self-analysis. *Journal of Management Studies* 35: 699–717.
- Carroll, J.S., and S. Hatakenaka. 2001. Driving organizational change in the midst of crisis. *MIT Sloan Management Review* 42: 70–79.
- Carroll, J.S., J.W. Rudolph, and S. Hatakenaka. 2002. Learning from experience in high-hazard organizations. *Research in Organizational Behavior* 24: 87–137.

- Carroll, J.S., J.W. Rudolph, S. Hatakenaka, T.L. Wiederhold, and M. Boldrini. 2001. Learning in the Context of Incident Investigation Team Diagnoses and Organizational Decisions at Four Nuclear Power Plants: Linking Expertise and Naturalistic Decision Making, E. Salas and G. Klein, eds. Mahwah, N.J.: Lawrence Erlbaum.
- Cook, S.D.N., and J.S. Brown. 1999. Bridging epistemologies: the generative dance between organizational knowledge and organizational knowing. *Organization Science* 10: 381–400.
- Crossan, M.M., H.W. Lane, and R.E. White. 1999. An organizational learning framework: from intuition to institution. *Academy of Management Review* 24: 522–537.
- deGeus, A. 1988. Planning as learning. *Harvard Business Review* 66(2): 70–74.
- Dougherty, D. 1992. Interpretive barriers to successful product innovation in large firms. *Organization Science* 3: 179–202.
- Gruenfeld, D.H., P.V. Martorana, and E.T. Fan. 2000. What do groups learn from their worldliest members: direct and indirect influence in dynamic teams. *Organizational Behavior and Human Decision Processes* 82: 45–59.
- Hansen, M.T. 1999. The search-transfer problem: the role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly* 44: 82–111.
- Hendershot, D.C., A.G. Keiter, J.W. Kacmar, P.C. Magee, P.C. Morton, and W. Duncan. 2003. Connections: how a pipe failure resulted in resizing vessel emergency relief systems. *Process Safety Progress* 22(1): 48–56.
- Jackson, S.A. 1996. Challenges for the Nuclear Power Industry and Its Regulators: The NRC Perspective. Speech presented at the Regulatory Information Conference, Washington, D.C., April 9, 1996.
- Jackson, S.E. 1992. Team Composition in Organizational Settings: Issues in Managing an Increasingly Diverse Workforce. Pp. 138–173 in *Group Process and Productivity*, S. Worshel, W. Wood, and J.A. Simpson, eds. Newbury Park, Calif.: Sage Publications.
- Jehn, K.A., G.B. Northcraft, and M.A. Neale. 1999. Why differences make a difference: a field study of diversity, conflict, and performance in workgroups. *Administrative Science Quarterly* 44: 741–763.
- Kemeny, J.G., B. Babbitt, P.E. Haggerty, C. Lewis, P.A. Marks, C.B. Marrett, L. McBride, H.C. McPherson, R.W. Peterson, T.H. Pigford, T.B. Taylor, and A.D. Trunk. 1979. Report of the President's Commission on the Accident at Three Mile Island. New York: Pergamon Press.
- Klein, G. 1998. *Source of Power: How People Make Decisions*. Cambridge, Mass.: MIT Press.
- Lanir, Z. 1986. *Fundamental Surprise*. Eugene, Ore.: Decision Research.
- March, J.G. 1991. Exploration and exploitation in organizational learning. *Organization Science* 2: 71–87.
- Marcus, A.A., P. Bromiley, and M. Nichols. 1989. *Organizational Learning in High Risk Technologies: Evidence from the Nuclear Power Industry*. Discussion Paper #138. Minneapolis: University of Minnesota Strategic Management Research Center.
- Marcus, A.A., and M.L. Nichols. 1999. On the edge: heeding the warnings of unusual events. *Organization Science* 10: 482–499.
- Nutt, P.C. 1999. Surprising but true: half the decisions in organizations fail. *Academy of Management Executive* 13: 75–90.
- Reason, J. 1990. *Human Error*. New York: Cambridge University Press.
- Reason, J. 1997. *Managing the Risks of Organizational Accidents*. Brookfield, Vt: Ashgate Publishers.
- Rochlin, G.I., and A. von Meier. 1994. Nuclear power operations: a cross-cultural perspective. *Annual Review of Energy and the Environment* 19: 153–187.
- Rudolph, J.W., E.G. Foldy, and S.S. Taylor. 2001. Collaborative Off-Line Reflection: A Way to Develop Skill in Action Science and Action Inquiry. Pp. 405–412 in *Handbook of Action Research*, P. Reason and H. Bradbury, eds. Thousand Oaks, Calif.: Sage Publications.

- Schulman, P.R. 1993. The negotiated order of organizational reliability. *Administration and Society* 25: 353–372.
- Senge, P. 1990. *The Fifth Discipline*. New York: Doubleday.
- Sitkin, S.B. 1992. Learning through failure: the strategy of small losses. *Research in Organizational Behavior* 14: 231–266.
- Szulanski, G. 1996. Exploring internal stickiness: impediments to the transfer of best practices within the firm. *Strategic Management Journal* 17: 27–43.
- van der Schaaf, T.W., D.A. Lucas, and A.R. Hale, eds. 1991. *Near Miss Reporting as a Safety Tool*. Oxford, U.K.: Butterworth-Heinemann.
- von Hippel, E. 1994. “Sticky information” and the locus of problem solving: implications for innovation. *Management Science* 40(4): 429–439.
- Weick, K.E., K.M. Sutcliffe, and D. Obstfeld. 1999. Organizing for high reliability: processes of collective mindfulness. *Research in Organizational Behavior* 21: 81–123.
- Weick, K.E., and K.M. Sutcliffe. 2001. *Managing the Unexpected: Assuring High Performance in an Age of Complexity*. San Francisco: Jossey-Bass.
- White, K.B. 1984. MIS project teams: an investigation of cognitive style implications. *MIS Quarterly* 8(2): 95–101.
- Wildavsky, A. 1988. *Searching for Safety*. New Brunswick, N.J.: Transaction Press.